PLACEMENT OF METADATA ON DATA STORAGE DRIVES IN A FIRST STORAGE ENCLOSURE OF A DATA STORAGE SYSTEM

TECHNICAL FIELD

[0001] The disclosed technology relates generally to storing metadata in a data storage system, and more specifically to technology for selectively storing data storage metadata on data storage drives contained in a first storage enclosure of a data storage system.

BACKGROUND

[0002] Data storage systems are arrangements of hardware and software that include non-volatile data storage devices (e.g. electronic flash drives, magnetic disk drives, and/or optical drives) that are used to provide non-volatile storage for processing I/O (Input/Output) requests (i.e. writes and/or reads) that are received from one or more hosts (e.g. host computers). The host I/O requests that are processed by a data storage system may include block I/O requests as well as file I/O requests. The host I/O requests indicate host data that is stored in logical volumes of non-volatile data storage that are contained in and maintained by the data storage system. The data storage system performs various data storage services that organize and secure host data received from host computers on the non-volatile data storage devices of the data storage system.

SUMMARY

[0003] While processing received host I/O requests and providing related data services, a data storage system generates and/or uses internal data storage metadata. For example, data storage metadata generated and/or used by a data storage system may include without limitation one or more of the following, as well as other specific types of metadata:

[0004] i) Mapping metadata (e.g. a mapping tree or the like) generated by and/or used by mapping logic to identify the locations of blocks of physical non-volatile storage that are mapped to corresponding portions of the logical address space of a logical volume accessed by host I/O requests received and processed by the data storage system,

[0005] ii) Mapped RAID (Redundant Array of Independent Disks) metadata generated by and/or used by mapped RAID logic, such as metadata (e.g. one or more tables, etc.) that describes how host data is striped across physical extents of non-volatile storage that are allocated to support mapped RAID,

[0006] iii) Snapshot metadata generated by and/or used by snapshot logic, such as metadata that indicates and/or describes (e.g. points to) point in time copies referred to as "snapshots" that the data storage system captures of storage objects such as logical volumes, LUNs (Logical Units), and/or other specific types of storage objects,

[0007] iv) Data deduplication metadata generated by and/or used by data deduplication logic, such as one or more data structures (e.g. a table or the like) that associate crypto-digests of previously stored blocks of host data with the locations in non-volatile storage at which the corresponding blocks of host data were stored,

[0008] v) Data compression metadata generated by and/or used by data compression logic, such as indications (e.g. pointers) of locations in non-volatile storage at which previously compressed host data is stored, and/or indications of the specific type(s) of compression (e.g. compression keys, compression algorithms, etc.) that were used to compress the previously compressed host data,

[0009] vi) Thin provisioning metadata generated by and/or used by thin provisioning logic, such as a data structure (e.g. a bit map) indicating which portions of the logical address space of a data storage object (e.g. a thin LUN or logical volume) have been allocated non-volatile data storage to handle received host I/O requests directed to the data storage object, and/or

[0010] vii) Logging metadata generated by and/or used by logging logic, such as a transaction log or the like storing metadata changes resulting from host I/O requests received by the data storage system, and that can be "replayed" to recover the data storage system to a consistent state after a system crash.

[0011] Access by the storage processor of a data storage system to internal data storage metadata is important for both the resiliency and the performance of a data storage system. For example, access to mapping metadata is critical for data storage system resiliency because loss of some or all of the mapping metadata may prevent the data storage system from locating host data that was previously stored in non-volatile data storage, resulting in data loss. The impact of losing even one block of data storage metadata can be very high. For example, loss of a single block of mapping metadata could potentially result in the loss of the mappings needed by the data storage system to access hundreds of gigabytes (GBs) of host data previously stored in non-volatile data storage.

[0012] Access by the storage processor to internal data storage metadata is also important for the performance of the data storage system with regard to processing host I/O requests. For example, in order to access host data previously stored in non-volatile data storage, multiple blocks of metadata may need to be read, e.g. in order to traverse a mapping tree that stores mappings between logical addresses of host data within a logical volume (e.g. offsets, logical block numbers, etc.) and corresponding physical locations of portions (e.g. blocks) of non-volatile data storage. As a result, improving access times for internal data storage metadata may improve access times for host data stored in non-volatile storage, and may have an even greater impact on performance of the data storage system with regard to I/O request processing than simply improving access times for the host data stored in non-volatile storage.

[0013] Various operational factors, such as the total amount of internal data storage metadata that must be stored, and/or other factors, may result in some amount of the internal data storage metadata generated and/or used by a data storage system being stored in the non-volatile data storage of the data storage system instead of memory. Some previous systems have exhibited shortcomings in this regard. In general, data storage systems may include multiple storage enclosures, each of which contains (e.g. has directly connected thereto) one or more non-volatile data storage devices. For purposes of explanation, the storage enclosures of a data storage system are referred to herein as a first storage enclosure that is directly connected to at least one storage processor of the data storage system, and one or more secondary storage enclosures that are indirectly connected to the storage processor. For example, a first storage